



STANFORD RESEARCH INSTITUTE
Menlo Park, California 94025 · U.S.A.

THE PROBABILISTIC STRUCTURE OF
PLANETARY CONTAMINATION MODELS

(NASA-CR-130558) THE PROBABILISTIC
STRUCTURE OF PLANETARY CONTAMINATION
MODELS (Stanford Research Inst.) - 33 p
HC \$3.75

N73-17103

CSCL 06M

G3/04

Unclas
62136

J. Michael Harrison

D. Warner North

Project Memo #1

SRI Project 2274

Presented at the Semiannual NASA
Spacecraft Sterilization Technology Seminar
New Orleans, Louisiana

January 30, 1973

I INTRODUCTION

A continuing source of concern in man's exploration of space is the potential for biological contamination of other planets. The possibility exists that a planet such as Mars is capable of supporting terrestrial microorganisms, in which case the landing of unsterilized space vehicles could lead to widespread contamination. One cause for alarm over this possibility is that it would greatly impede subsequent attempts to detect indigenous life. In addition, many experts simply feel that contamination would be morally reprehensible, and it has also been suggested that a proliferation of earthly microbes might impede future efforts toward colonization.¹ The apparent solution to these concerns is vigorous sterilization of all spacecraft engaged in planetary exploration. Sterilization can be costly, however, in terms of both direct expense and reduced spacecraft reliability.

In response to concern over potential contamination of Mars, a planetary quarantine resolution was adopted in 1966 by the Committee on Space Research (COSPAR) of the International Council of Scientific Unions. This resolution required that the spacefaring nations conduct their unmanned explorations of Mars in such a way that the total probability of contamination during a specified quarantine period not exceed 10^{-3} . In its desire to conform with this standard, the National Aeronautics and Space Administration (NASA) has established a planetary quarantine policy, administered by a planetary quarantine officer (PQO) whose responsibilities are broadly depicted in Figure 1. Given tentative plans as to the number and character of future missions to Mars, the quarantine officer establishes a maximum permissible value for the probability of contamination

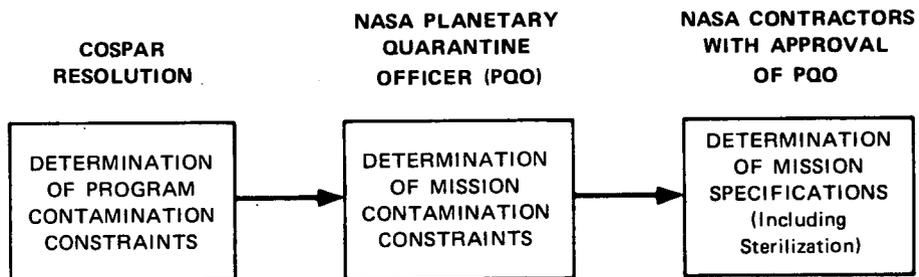


FIGURE 1 THE HEIRARCHY OF PLANETARY QUARANTINE DECISIONS

from individual missions. These upper bounds, widely known as mission "allocations," are set in such a way as to ensure program conformance with the COSPAR constraint. The final level of decision-making is then the determination of engineering specifications that will most efficiently meet the mission constraint. Some primary means of reducing contamination probabilities from individual missions are trajectory biasing for fly-bys, altitude restrictions for orbiters, and spacecraft sterilization for landers. The central analytical problem of the planetary quarantine program is to determine the probability of contamination under a variety of different mission specifications. This we call the "mission contamination problem." Its solution is a prerequisite for the determination of which specifications meet a given mission constraint.

The purpose of this paper is to discuss at a broad conceptual level the analytical basis for planetary quarantine standards and procedures. Attention is focused on the final phase of the decision process shown in Figure 1, and the discussion centers around only the case of landing

missions. To a large extent, current standards and procedures are based on the probabilistic model of planetary contamination advanced by Sagan and Coleman,² whose work provided both the stimulus and the theoretical foundation for the COSPAR resolution. We begin in Section II with a brief account of the Sagan-Coleman model, pointing out the implications of certain independence assumptions that are critical to their analysis. It is our contention that these assumptions give an overly simplified characterization of the problem and that their relaxation might lead to substantial changes in the final probability assessments. Thus the issue is not just one of theoretical interest, but one whose potential implications for policy are considerable. It should be emphasized that our criticisms concern the structure of the Sagan-Coleman model rather than the numerical values assigned to various input parameters. On the matter of parameter values, there have been considerable debate and discussion^{3,4} but the basic structural assumptions and resulting formulas are widely accepted by COSPAR, NASA, and NASA contractors as a means of determining sterilization requirements for Project Viking and other future unmanned planetary missions.

In Section III we discuss a classical problem in probability theory that provides, at least in our opinion, a close conceptual parallel to the type of dependence present in the contamination problem. It is this dependence that the Sagan-Coleman model fails to account for. In Section IV we indicate how the shortcomings of the Sagan-Coleman analysis can be remedied by adopting a slightly richer model structure, explicitly modeling more of the contingencies that underlie our uncertainty about contamination.

A general theme underlying our entire presentation is the firm conviction that any analytical treatment of planetary quarantine requirements must lie squarely in the domain of subjective probability assessment.

It is true that much physical evidence exists relevant to the contamination issue (such as the findings of Mariner 9) and that large amounts of data (such as the lethality of various sterilization treatments) are available on various aspects of the problem. It is inevitable, however, that the quantification of expert scientific judgment will provide the ultimate means through which all available information is integrated into an assessment of the total risk of contamination. The model builder's task is to structure a problem so as to ease the burden of direct assessment as much as possible without imposing or presuming relationships that may conflict with the expert's judgment. The weaknesses of the Sagan-Coleman model do not arise from any erroneous characterization of physical processes or objects; the question is whether the assumptions of the model are consistent with the current state of scientific information. A separate issue is the degree to which human beings, even very sophisticated ones, are capable of directly assessing probabilities as small as 10^{-4} , which the Sagan-Coleman model requires experts to do. In Section V we indicate how an enriched model structure can help to overcome this difficulty and also to facilitate communication among experts. Finally, in Section VI the role of subjective probability assessments in scientific investigation is discussed briefly, with particular emphasis on the meaning of the COSPAR resolution.

II THE SAGAN-COLEMAN FORMULA

Let us consider a single mission to Mars. To be specific, we will suppose that the mission is intended to land a vehicle at some specified equatorial site on Mars some time in July 1976. Given all mission specifications, including a proposed sterilization procedure, the mission contamination problem is to determine the probability of

C = the event that Mars will be biologically contaminated by organisms aboard this spacecraft

in terms of more fundamental descriptors of the mission. Toward that end, we further define

N = the number of viable organisms released to the Martian environment or into its atmosphere from the spacecraft (a random variable),

$$m = \sum_{k=1}^{\infty} k \text{ Prob}\{N=k\}$$

= the expected (or mean) number of viable organisms released,

P_g = the probability that a single released organism will survive, multiply, and contaminate a significant fraction of the planet.

The solution advanced by Sagan and Coleman for the mission contamination problem was the simple approximation

$$\text{Prob}\{C\} \approx m P_g, \quad (1)$$

for which no justification was offered. The apparent rationale for this approximation is as follows. Defining

E_i = the event that the i^{th} released organism does not survive to multiply and cause contamination,

it is immediate that

$$1 - \text{Prob}\{C\} = \text{Prob}\{E_1 \text{ and } \dots \text{ and } E_N\}$$

$$= \sum_{k=1}^{\infty} \text{Prob}\{N=k\} \text{Prob}\{E_1 \text{ and } \dots \text{ and } E_k | N=k\} \quad (2)$$

If, given that k organisms are released, we assume the events E_1, E_2, \dots, E_k to be independent* and of equal probability, then

$$\text{Prob}\{E_1 \text{ and } \dots \text{ and } E_k | N=k\} = [\text{Prob}\{E_1 | N=k\}]^k \quad (3)$$

Moreover, if we assume that the survival of any one organism is independent of the number of organisms released, then

$$\text{Prob}\{E_1 | N=k\} = \text{Prob}\{E_1\} = 1 - P_g \quad (4)$$

Substituting (3) and (4) into (2) then gives us

$$1 - \text{Prob}\{C\} = \sum_{k=1}^{\infty} \text{Prob}\{N=k\} (1 - P_g)^k$$

$$\approx \sum_{k=1}^{\infty} \text{Prob}\{N=k\} (1 - k P_g) = 1 - m P_g \quad ,$$

* A set of events E_1, E_2, \dots, E_k is said to be independent if

$$\text{Prob}\{E_1 \text{ and } \dots \text{ and } E_i\} = \text{Prob}\{E_1\} \dots \text{Prob}\{E_i\}$$

for each $i = 1, 2, \dots, k$.

the approximation being justified by the fact that P_g is very small, and only values of N that are small compared with $1/P_g$ have significant probability.* This relationship is equivalent to (1).

Let us now consider the independence assumption underlying (3). If the assumption is accepted, then the following must hold by definition. Having learned that exactly k organisms were released and that none of the first $k-1$ survived, we would not be inclined to alter our initial assessment for the probability that the last will survive and proliferate. We do not think this is a reasonable statement, and the reason lies in a rather fundamental question. Why are we uncertain about the survival of any single organism? To a large extent, it is because many important characteristics of the Martian environment are not yet known with certainty. The events E_1, \dots, E_k are mutually dependent on the actual character of that environment. Learning the fate of the first $k-1$ organisms tells us something about Mars itself, which in turn tells us something about the k^{th} organism's chances of surviving and proliferating.

The extreme case of total dependence is illustrated by the following scenario. Suppose that the current state of scientific knowledge admits only two possibilities with respect to the Martian environment: either it is so hostile that no terrestrial organism could possibly survive, or else it is so hospitable as to ensure the survival and proliferation of any such organism. In this case the events E_1, E_2, \dots are totally dependent, since knowing the fate of any one organism would automatically tell us the fate of the others. If we assign a probability of 10^{-4} to the hospitable situation and $1-10^{-4}$ to the hostile one, then $P_g = 10^{-4}$. Note, however, that

* Under these conditions, the second and higher order terms in the expansion of $(1-P_g)^k$ are negligible.

$$\text{Prob}\{C\} = \text{Prob}\{N \geq 1\} \times P_g = \text{Prob}\{N \geq 1\} \times (10^{-4}) \quad (5)$$

In general, the value for $\text{Prob}\{C\}$ given by (5) will be much smaller than that given by (1). The actual state of current scientific information would put us somewhere between the case of independence, on which (1) is based, and the case of total dependence, on which (5) is based.

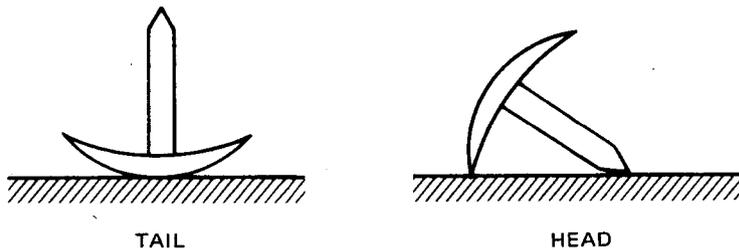
Turning to the independence assumption underlying (4), our objection is more subtle. Imagine that, as our spacecraft arrives on Mars, a scientific expert is sequestered and denied any information other than unimpeachable evidence that exactly ten viable organisms have been released. He is then asked to assess the probability that the first organism released will survive and proliferate. Would this assessment be any different if he were told that 1000 organisms had been released? Although we are not certain, there seems to be at least one plausible reason why he might give different answers in the two circumstances. Having learned that 1000 organisms were released, the scientist might think it likely that some engineering failure has led to a hard landing, releasing many encapsulated organisms. This in turn increases the likelihood that the vehicle has impacted far from its targeted landing site. A critical uncertainty regarding the life-supporting capability of Mars concerns the existence of liquid water on the planet, and most scientists feel that if liquid water exists at all, it is most likely to be in the polar regions. Consequently, knowing that many organisms have been released might alter the expert's assessment as to the accessibility of liquid water and hence the survivability of an individual organism. In summary, the number of released organisms depends on the mode of landing, which in turn affects survivability. This argues that the events $\{N=k\}$ and E_1 are not independent and hence that (4) is invalid. Another way of looking at this dependence is to say that Mars is not in fact characterized by a single P_g value. The probability that a single released

organism will survive and multiply depends on where the organism is released. Consequently, it is difficult for an expert to make meaningful probability assessments within the framework of a model that does not take explicit account of uncertainty about the actual landing site and impact velocity. This point is pursued further below.

The Sagan-Coleman linear approximation $\text{Prob}\{C\} \approx m P_g$ is a cornerstone of current quarantine planning procedures. To illustrate its use, let us suppose that the mission allocation requires $\text{Prob}\{C\} \leq 10^{-4}$. Then the corresponding constraint on the mean number of viable organisms released is $m \leq 10^{-4} / P_g$. The probability P_g , which is obviously very small, has been directly assessed by scientific experts, internally integrating the many factors that affect an organism's survivability on Mars. In contrast, a great deal of laboratory experimentation and additional modeling has been used in developing assessments for m under various sterilization procedures.⁵ We have argued that the linear approximation is based on independence assumptions that are highly suspect. In the following sections, we discuss the general notion of dependent events in a more familiar setting and then indicate how the Sagan-Coleman formulation can be modified to eliminate its current weaknesses.

III A CLASSICAL PARALLEL PROBLEM

The following problem, which has been discussed in a slightly different form by Howard,⁶ provides an example of a familiar physical process having identical but informationally dependent trials. Its significance to the mission contamination problem will be discussed shortly. Let us suppose that a tack is dropped onto a large flat surface, its two possible landing positions being labelled "heads" and "tails" as in the following diagram. You are told only that the tack in the diagram



is drawn to scale, that it will be dropped from a height of four feet by a human being, and that the surface is very flat. Your first problem is to assess the probability of a head in one toss, and your second is to assess the probability of ten heads in ten tosses. To respond that you don't know the probabilities, having never watched any tack tossing, is unacceptable. The questions do not concern frequencies or any other type of "physical fact." We ask only a quantification of your judgment, recognizing that different people will typically make different assessments. Now suppose that after much scrutiny of the diagram you assess

the probability of a head in one toss to be one-half. Is it possible to deduce from this response, using only the rules of consistency imposed by probability theory, your response to the second question? The answer is no. You simply have not told us enough about your judgment (or state of information). Before any calculations can be done (on your behalf), we need to know something about how you feel the individual tosses relate to one another. To fill in this gap, we might assume that you view the events,

B_i = the event of a head on the i^{th} toss ($1 \leq i \leq 10$) ,

as independent, in which case we immediately have

$$\text{Prob}\{\text{all heads in 10 tosses}\} = (1/2)^{10} \approx 10^{-3} .$$

But, considering the characterization of independence given earlier, does this assumption accurately reflect your state of information? It seems unlikely, for undoubtedly you would be inclined to alter your initial assessment for the probability of a head in one toss if we told you the results of the first nine tosses. Having rejected the independence assumption, how can you compactly express the degree of dependence that you perceive to exist among the results of the separate trials? Under very mild assumptions,* it can be shown that the following characterization provides all the required information. Let

ϕ = the fraction of heads that would be observed in a very long sequence of tosses,

which can be viewed only as a random variable with your current state of information. What we need is your subjective (prior) probability distribution for the random variable ϕ . This is conveniently expressed by the cumulative distribution function

* The assumption is that the trials be exchangeable. For a definition and discussion of exchangeable trials see de Finetti."

$$F(x) = \text{Prob}\{\phi \leq x\} \quad , \quad 0 \leq x \leq 1$$

The mean (or expected value) of this distribution is given by*

$$E(\phi) = \int_0^1 [1-F(x)] dx \quad ,$$

and consistency demands that it equal 0.5. That is, the axioms of probability theory require that your subjective probability of a head in one trial equal the mean of your subjective distribution for the fraction of heads in a great many trials.

Figures 2 through 4 show three possible distribution functions for the random variable ϕ , each of which is consistent with the earlier assessment that $\text{Prob}\{\text{head}\} = \text{one-half}$. The first of these distributions corresponds to the case of independent trials, the subject being absolutely certain that the long-run fraction of heads will be 50 percent. Such a distribution might be assessed by an individual who has spent the last few months tossing this same tack onto this same surface. Although he is uncertain as to what will happen in a few trials, he has complete knowledge of the basic environment, knowing that the tack is equivalent to an unbiased coin.

The second distribution (Figure 3) corresponds to the case of totally dependent trials. The subject is absolutely certain that the tack will always either come up heads or come up tails, but he is not sure which of these cases pertains. (He might have an acquaintance who has tossed the tack many times, told him it always falls one way, but left him to guess from there.) He has assessed the probability of all heads

* Integration by parts shows this formula equivalent to the usual one in terms of the density function or probability mass function.

$$F_1(x) = \text{PROB} \{ \phi \leq x \}$$

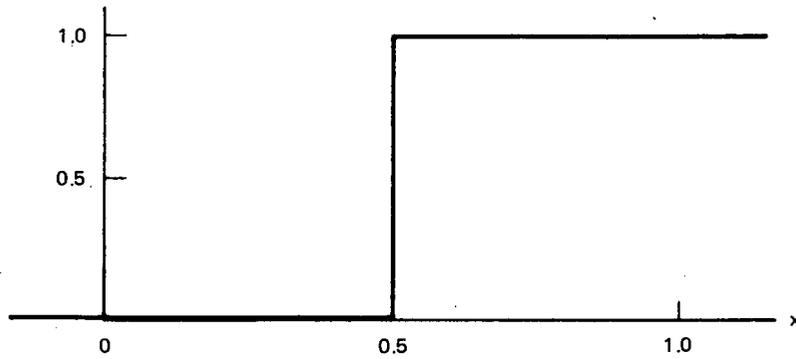


FIGURE 2 DISTRIBUTION IMPLYING INDEPENDENT TRIALS

$$F_2(x) = \text{PROB} \{ \phi \leq x \}$$

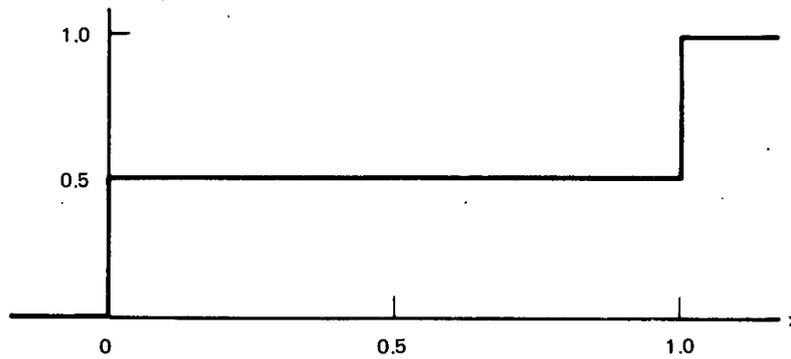


FIGURE 3 DISTRIBUTION IMPLYING PERFECTLY DEPENDENT TRIALS

$$F_3(x) = \text{PROB} \{ \phi \leq x \}$$

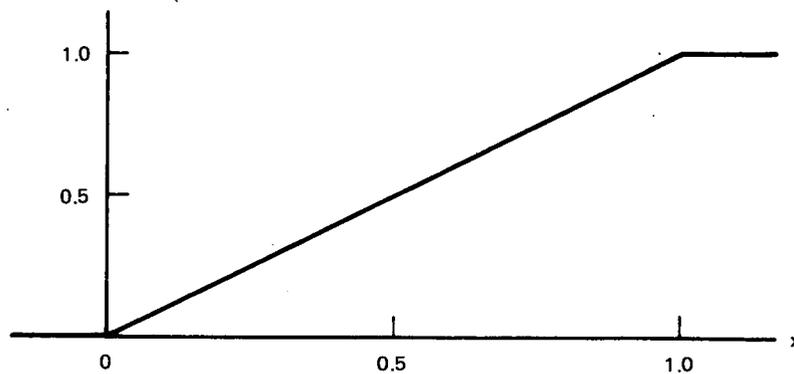


FIGURE 4 UNIFORM DISTRIBUTION

3

to be one-half and that of all tails to be one-half. Note that if this subject were able to observe one toss it would resolve all his uncertainty regarding the outcomes of subsequent tosses.

The type of distribution that we would generally expect, intermediate to the preceding extreme cases, is shown in Figure 4. Here the subject reveals great uncertainty as to the experiment's environment, assigning a uniform distribution over the interval of possible values. The mean of his distribution, like that of the others, is $E(\phi) = \text{one-half}$.

Given the probability distribution for ϕ , we can calculate the probability of all heads in n trials using the formula*

$$\text{Prob}\{\text{all heads in } n \text{ trials}\} = E(\phi^n), n \geq 1 .$$

From this we have computed the relationships shown in Figure 5 for each distribution discussed earlier. The subject who views the trials as independent thinks it very unlikely (less than one chance in a thousand) that we could survive the ten trials without observing a tail. In contrast, the subject who views the trials as perfectly dependent continues to assign a probability of one-half to the event of all heads, regardless of how many times the tack is tossed. The corresponding relationship for the third subject lies between these two extremes. In particular, he assesses the probability of ten heads in ten trials to be about 9 percent, one hundred times the probability assigned by the second distribution. Thus we find that the three individuals differ greatly in their assessment of what is likely to occur in repeated trials, although they agree perfectly as to the probability of a head in a single trial. It is the degree of informational dependence among trials that differs from one subject to another, and these differences have significant implications.

* This is an application of de Finetti's theorem.⁷

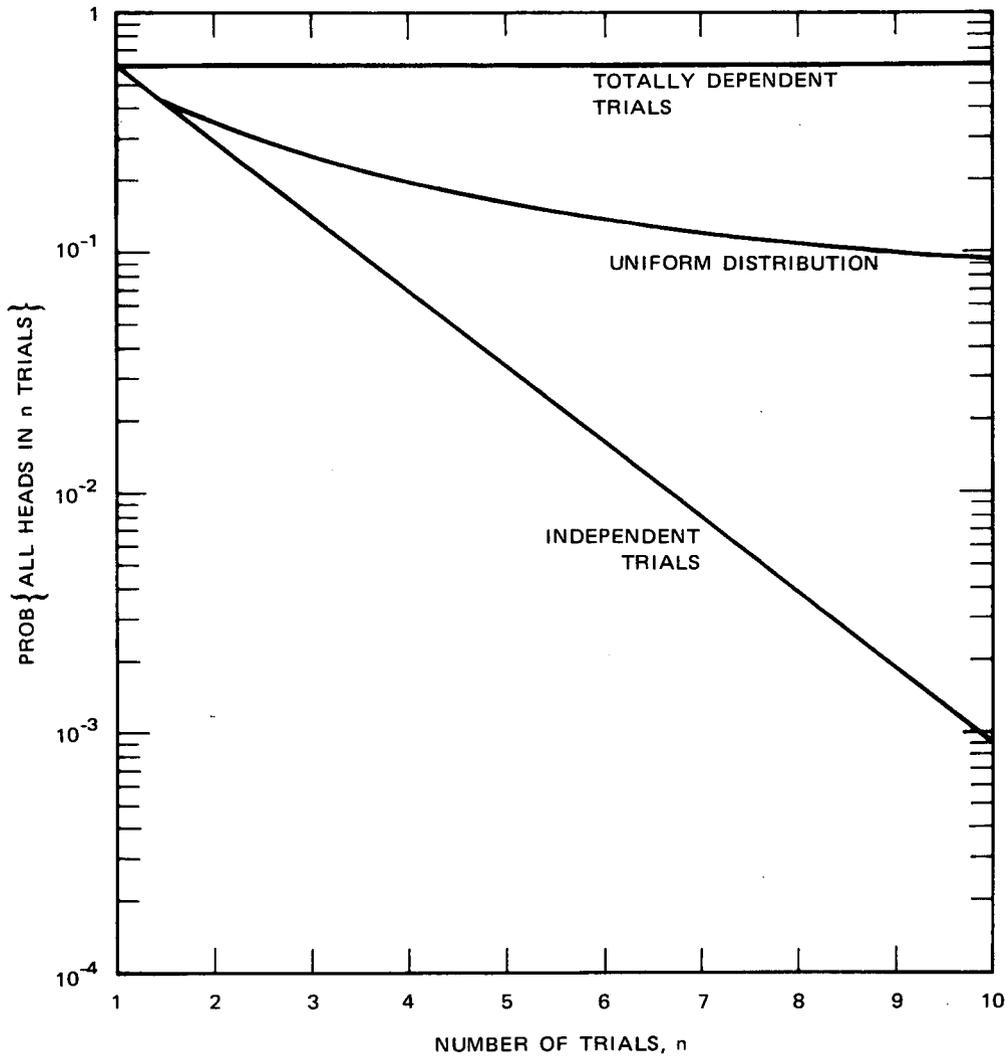


FIGURE 5 PROBABILITY OF ALL HEADS VERSUS NUMBER OF TRIALS

Returning to our original concern, consider the following restricted contamination problem. Suppose that our unmanned mission is a complete success, landing softly at the targeted equatorial site. We shall denote this event of successful landing by S . The problem is to assess $\text{Prob}\{C|S, N=k\}$, the probability of contamination given that the landing is successful

and k viable organisms are released, as a function of k . The linear approximation

$$\text{Prob}\{C|S, N=k\} \approx k P_g$$

is represented by the upper curve in Figure 6, using an illustrative P_g value of 10^{-4} . To draw a parallel between our thumbtack example and the problem at hand, let us associate the fate of each individual organism with a separate toss of the tack, the event of survival and proliferation being associated with the outcome "tails" and the organism's demise being associated with "heads." Then $1 - \text{Prob}\{C|S, N=k\}$ corresponds to the probability of all heads in k trials. The Sagan-Coleman linear approximation is based on an assumption of independent trials, viewing the individual organisms as severely biased tacks whose degree of bias is precisely known. In marked contrast, the extreme scenario of perfect

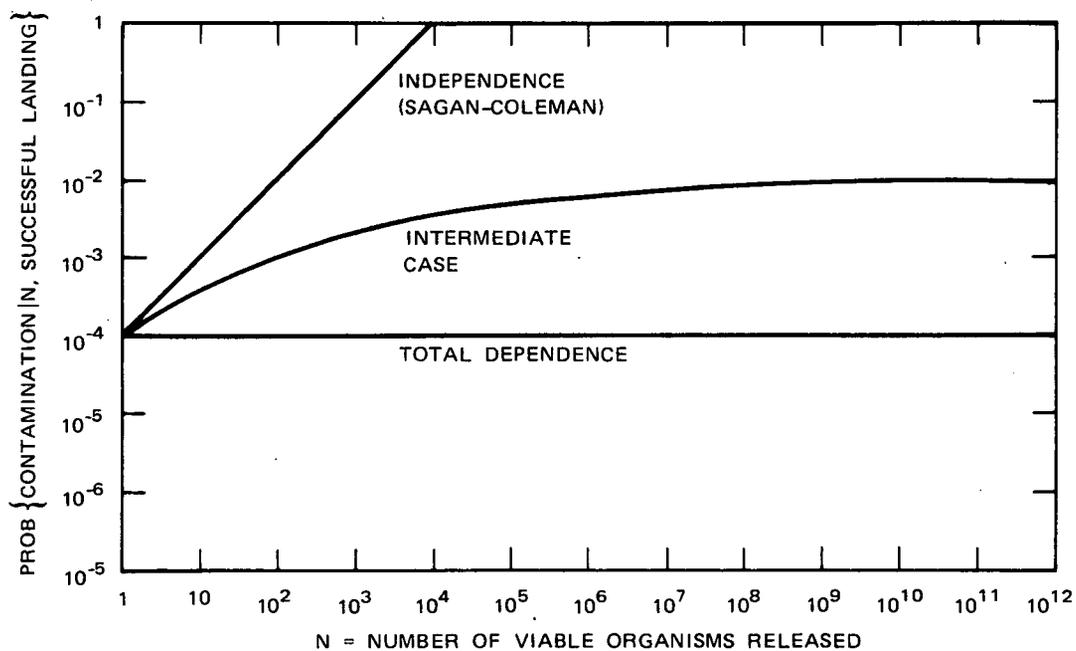


FIGURE 6 ILLUSTRATIVE SOLUTIONS FOR THE RESTRICTED CONTAMINATION PROBLEM

dependence is represented by the lower curve in Figure 6. Here the subject is certain that either all the organisms will survive or none will, so $\text{Prob}\{C|S, N=k\}$ is a constant for all $k \leq 1$. An intermediate case is represented by the middle curve in Figure 6. This subject, like both the others, assesses the probability of a single organisms's survival to be 10^{-4} . He feels that $\text{Prob}\{C|S, N=k\}$ increases with k , but not in direct proportion, approaching an asymptotic value of 10^{-2} as k gets very large. Thus 10^{-2} represents the subject's assessed probability that Mars is contaminable.

The final "solution" to the restricted contamination problem is a curve like those shown in Figure 6. One could conceivably compute this relationship by analogy to the thumbtack example, first assessing a distribution for ϕ , the fraction of a great many organisms that would die in a hypothetical experiment, but we feel that better methods are available. What the thumbtack analogy illustrates is the basic artificiality of the independence assumption behind (3) and the potential impact of accounting computationally for informational dependence.

IV ENRICHING THE MODEL STRUCTURE

Returning to the overall mission contamination problem, let us consider the various levels of detail at which this problem can be modeled. The first possibility is to do no modeling whatsoever, simply telling the expert our mission specifications and asking him to assess $\text{Prob}\{C\}$ directly. The level of detail proposed by Sagan and Coleman is represented in Figure 7. Here the basic structure consists of a "biorelease model," whose output is the mean number of released organisms, and a "proliferation model," consisting of a linear relationship between number of released organisms and contamination probability. It is our feeling that the Sagan-Coleman model structure must be enriched to at least the extent shown in Figure 8 if assessment is to be meaningful. An initial "landing model" has been added through which we explicitly express our uncertainty about the technical success of the mission. As Figure 9 indicates, the role of the landing model is to identify a relatively small number of mutually exclusive and exhaustive possibilities with respect to landing

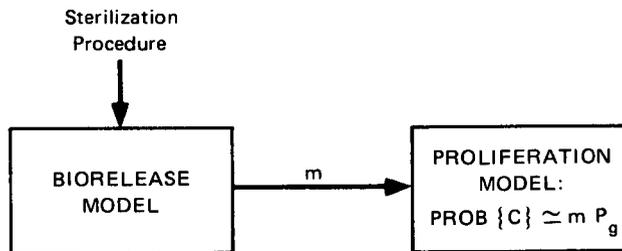


FIGURE 7 BASIC STRUCTURE OF THE SAGAN-COLEMAN MODEL

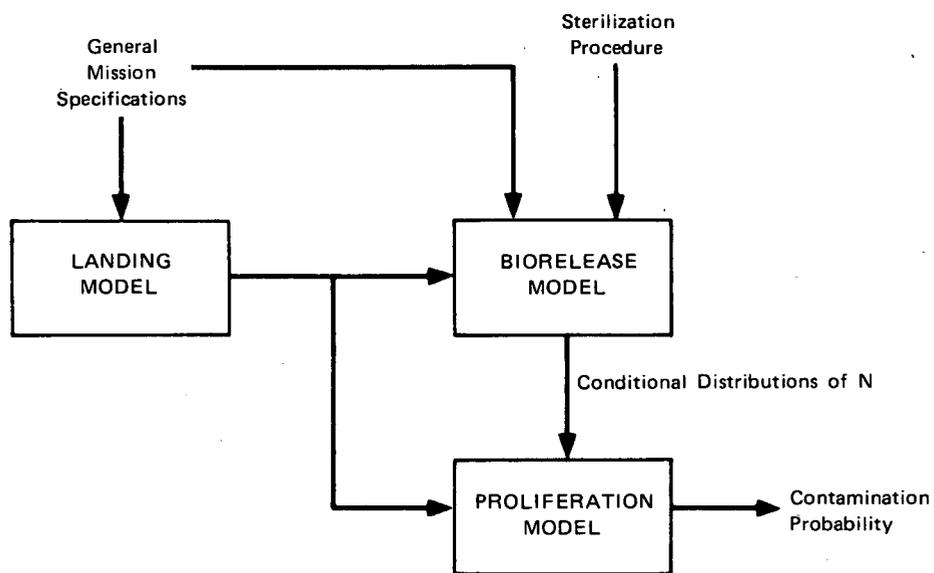


FIGURE 8 BASIC LOGICAL STRUCTURE OF A MISSION CONTAMINATION MODEL

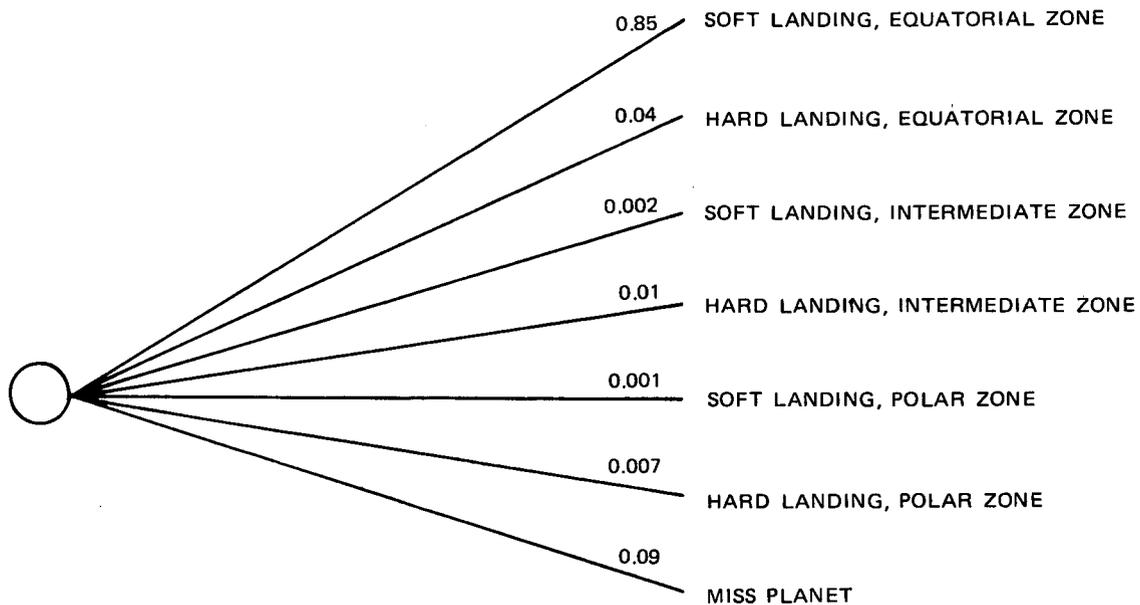


FIGURE 9 ILLUSTRATIVE OUTPUT OF LANDING MODEL (ASSUMING TARGETED LANDING SITE IS EQUATORIAL)

site and impact velocity, associating a probability assessment with each such possibility. The other major model components are then a biorelease model, yielding a probability distribution for N that is conditional on the landing outcome, and a proliferation model, yielding a functional relationship between $\text{Prob}\{C\}$ and N that is also conditional on the landing outcome. Figure 10 presents an illustrative output of the biorelease model conditional on a specific landing outcome. A probability is associated with each of several possible ranges of values that N might assume. A similar illustrative output of the proliferation model, also conditional on a specific landing outcome, is shown in Figure 11. We have shown the functional relationship between $\text{Prob}\{C\}$ and N to be

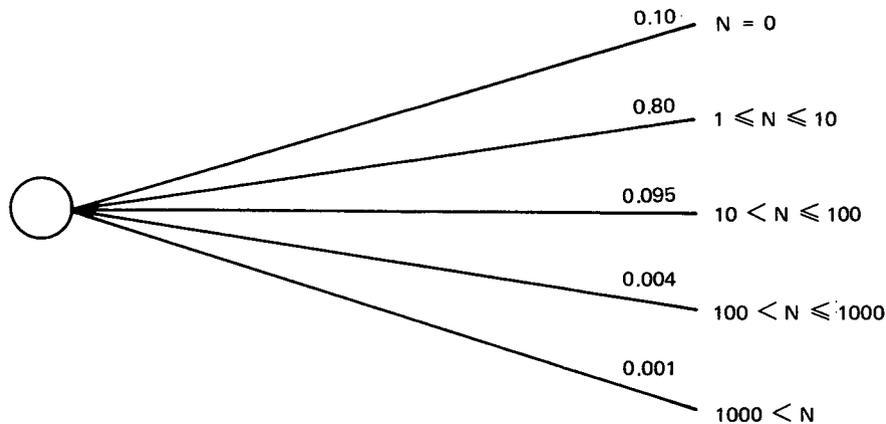


FIGURE 10 ILLUSTRATIVE OUTPUT OF BIORELEASE MODEL (GIVEN SOFT EQUATORIAL LANDING)

nonlinear, although the numerical values are purely hypothetical, to emphasize again that one cannot assume the survival of individual organisms to be independent events.

The purpose of the landing model in our proposed logical structure is to provide enough information about technical success of the mission that the number of released organisms and the survival of any one organism become independent given this information. It is this requirement that must dictate the level of detail at which landing outcomes are specified. When all three major model components have been built, the total probability of contamination is computed in the obvious way, multiplying and adding according to the laws of conditional probability. In this regard one point is worth noting. The Sagan-Coleman analysis implies that the

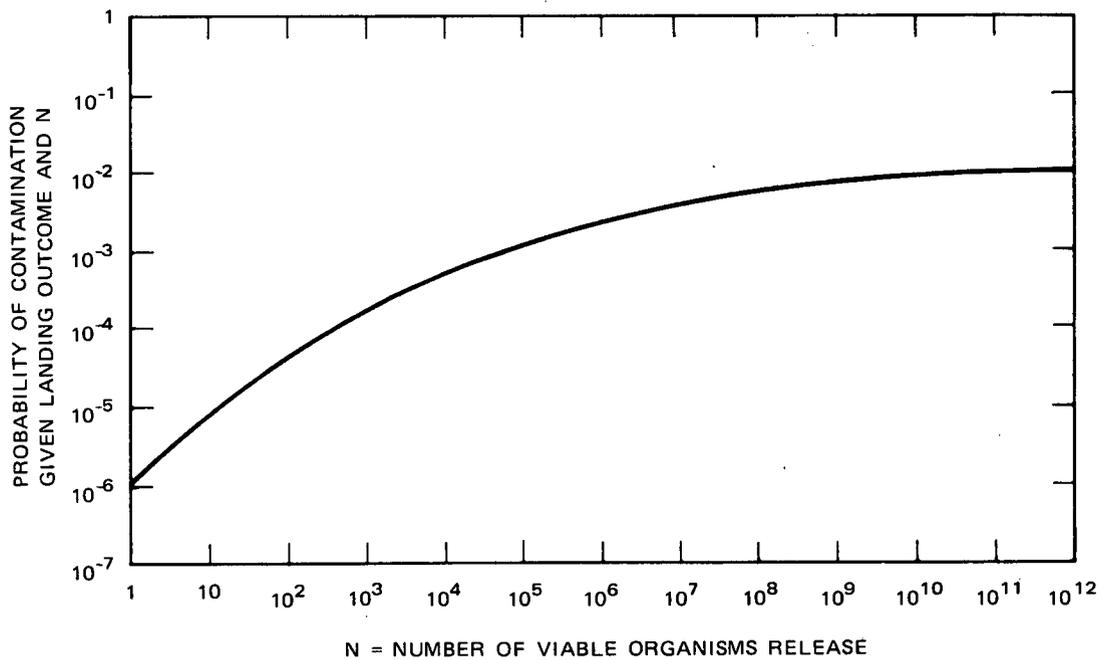


FIGURE 11 ILLUSTRATIVE OUTPUT OF PROLIFERATION MODEL (GIVEN SOFT EQUATORIAL LANDING)

only information about the distribution of N that we need to know in computing $\text{Prob}\{C\}$ is its mean. As the computations in Section II demonstrate, this results directly from the assumed linear relationship between $\text{Prob}\{C\}$ and N . If we abandon the Sagan-Coleman independence assumption, recognizing that for any given landing outcome the relationship may be nonlinear, then the computed value of $\text{Prob}\{C\}$ will depend in general on the entire distribution of N rather than just its mean.

The reader will note that the illustrative proliferation model output in Figure 11 serves to solve what we earlier called a restricted contamination problem for a specified landing outcome. In Section III it was argued that the thumbtack problem contains a similar type of

uncertainty, but we do not feel that a similar method of solution is indicated for the two problems. If someone were to ask exactly why you view successive tosses of a thumbtack as dependent trials, it would probably be hard to answer in physical terms. The uncertainty that we feel about the outcomes of individual trials is difficult to express directly in terms of the uncertainty surrounding such physical properties as the hardness of the surface, the weight of the tack head, and so forth. Rather it is the extremely complicated interaction of many considerations, most of them not consciously identified, that leads to uncertainty. For most people, the best way to think about the problem is through the long-run frequency ϕ identified earlier. Equivalently stated, it is difficult (or unnatural, or unprofitable) to model the problem at a level of detail richer than a single probability distribution for ϕ .

In contrast, most experts find it quite easy to identify specific physical considerations that contribute to their uncertainty about an organism's survival on Mars. Some of the most important such considerations are schematically represented in Figure 12. The central line in this figure portrays basic events required for growth and proliferation of a released organism, while the upper and lower lines portray characteristics of the Martian environment that influence the likelihood of those events. Of critical importance is the matter of liquid water. There are no known terrestrial organisms that can proliferate in the absence of water, and it is thought that the atmospheric pressure everywhere on Mars is below the triple-point pressure for water. Thus an organism must find its way to a nonequilibrium microenvironment, such as a salt pool or a pressure pocket under a polar cap, where water exists at least periodically. Important environmental factors are then the existence, general location, extent, and accessibility of such hospitable microenvironments. Also important are the degree of protection from

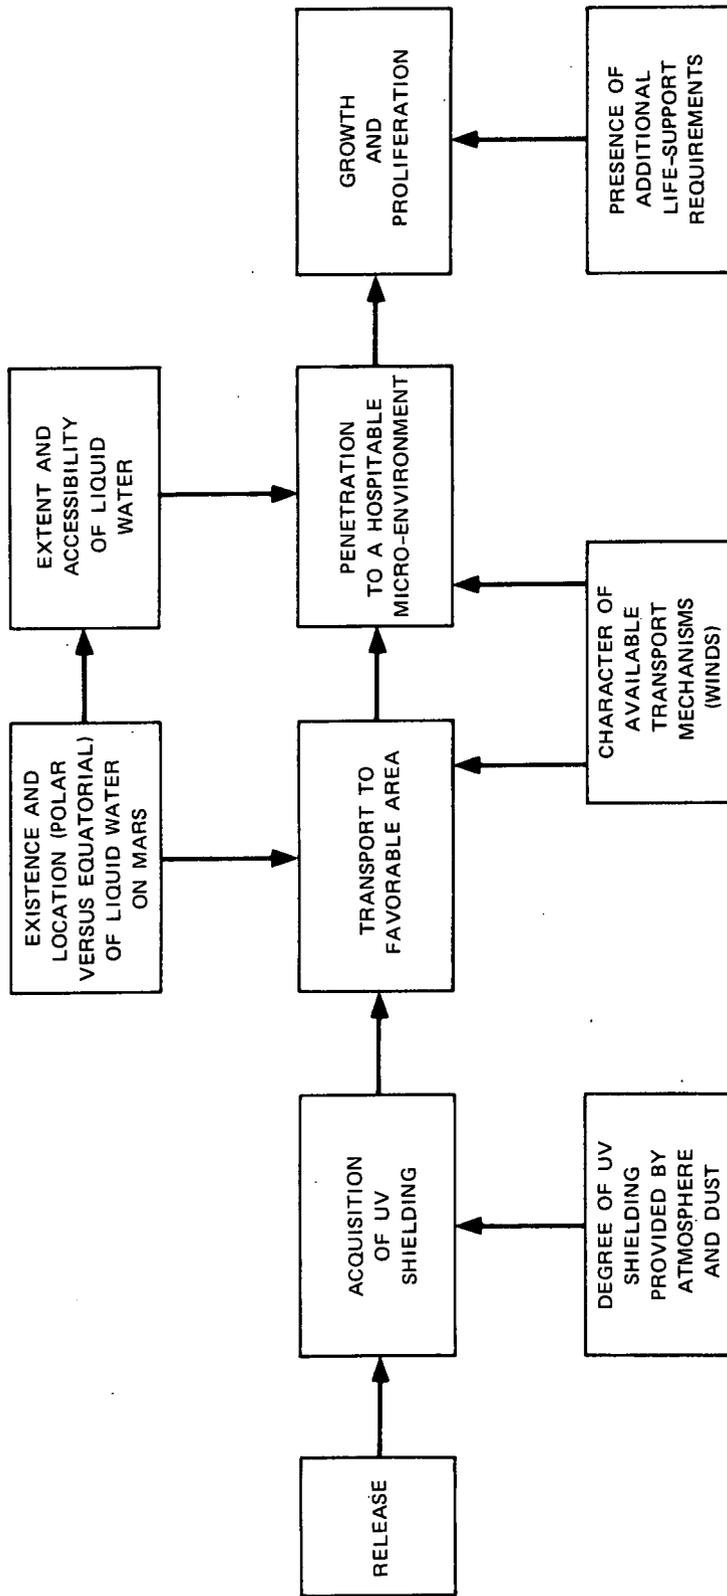


FIGURE 12 EVENTS AND PROCESSES CRITICAL TO BIOLOGICAL CONTAMINATION

ultraviolet radiation afforded by the Martian atmosphere and dust, which bears on the organism's survival in transit, and the character of Martian wind patterns. Even if we were given a relatively complete description of all these environmental factors, a good deal of uncertainty would remain concerning the organism's fate, and the events on the central line of Figure 12 are intended to reflect that remaining uncertainty.

It is conceivable that each of the functional relationships between $\text{Prob}\{C\}$ and N that constitute the "proliferation model" shown in Figure 7 could be directly encoded from expert judgment. This would require that the expert integrate internally all the diverse considerations shown in Figure 12. At the other extreme, one could build an exhaustive probabilistic model of the Martian environment and then encode the relationship conditional on both the landing outcome and the values assumed by all environmental factors. It is conceivable that we could condition on so many environmental factors as to make the survival of individual organisms independent given this information, in which case an expert would have to assess only the conditional probability of a single organism's survival and proliferation for each combination of environmental factors and landing outcome. Between these two extremes of direct encoding and exhaustive modeling of the Martian environment are intermediate levels of detail similar to those shown in Figure 13. Here we explicitly represent uncertainty about a few of the environmental factors that make the survival of individual organisms mutually dependent. The relationship between $\text{Prob}\{C\}$ and N is then directly encoded for each combination of values for the factors modeled, always conditional on a specific landing outcome. The underlying model structure shown in Figure 13 is very crude, and it is all but impossible to say what represents an "appropriate" level of detail within the broad range of potential models. Typically, one should begin with a very crude structure and continue refining it so long as experts feel that the additional detail is helpful in focusing their judgment.

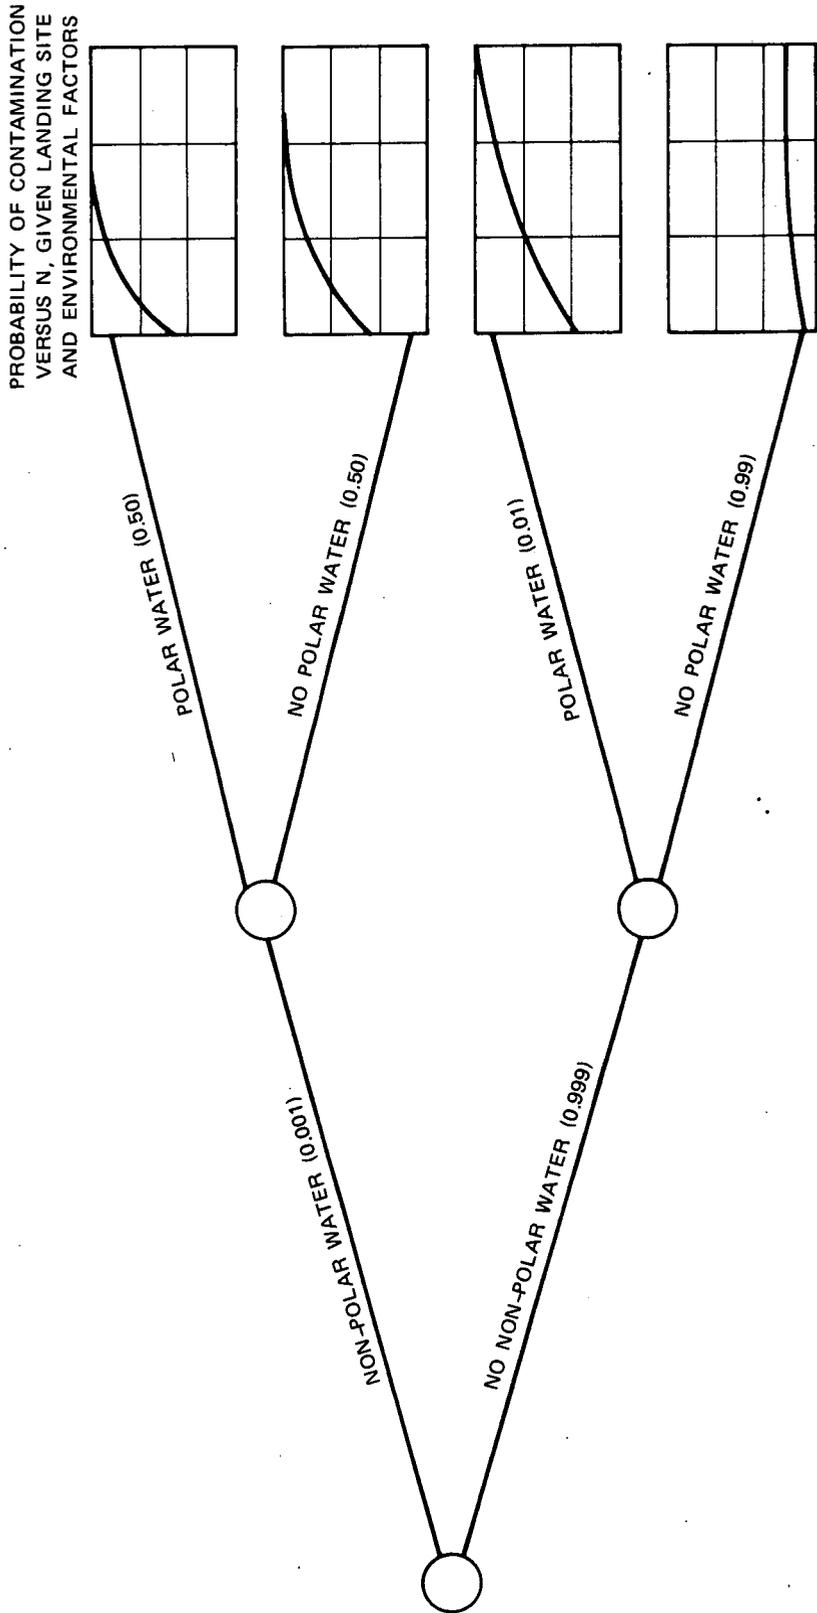


FIGURE 13 ILLUSTRATIVE DETAIL FOR PROLIFERATION MODEL (GIVEN SOFT EQUATORIAL LANDING)

V THE VALUE OF MODELING FOR ENCODING AND COMMUNICATION

We have indicated that a richer model structure would be valuable for the problem of assessing $\text{Prob}\{C\}$ in terms of N (given a specified landing site). There are two basic reasons for this recommendation. First is that more detailed modeling will help the expert in assessing a single organism's chances for survival and proliferation. That is, he can build up a personal assessment for P_g by first making separate assessments on various critical environmental factors and then assessing P_g conditional on these factors. This accomplishes a reduction in complexity, replacing one large, complicated question with a series of smaller ones that are more easily conceptualized. We have also argued that a richer model structure will assist the expert in assessing the parametric dependence of $\text{Prob}\{C\}$ on N by explicitly accounting for at least some of the factors that make the survival of separate organisms informationally dependent events.

Two other advantages to additional modeling relate closely to the reduction of complexity. The first concerns the well-documented difficulty of directly assessing very small probabilities. If a person tells you that he assesses the probability of event E_1 to be 10^{-4} and that of event E_2 to be 10^{-6} , then you can be quite sure that he considers both events unlikely, with E_1 more likely than E_2 . Experience indicates, however, that we should take care in attaching any absolute significance to the numerical assessments. Would he, for example, rather bet on E_2 occurring or on being dealt a royal flush in a game of 5-card stud poker? Since calculation will show that the latter event has a probability of about 1.5×10^{-6} , the person's assessment would lead us to conclude that he prefers to bet on the royal flush. Yet one can hardly be confident

that such a preference would emerge if the question were asked directly. The problem is that, when asked to assess probabilities smaller than (say) 1/100, we all have difficulty conjuring up familiar reference events that we perceive to be of comparable likelihood. In many applications, a probability of 10^{-2} or 10^{-3} can in fact be used as a working definition of impossibility. One might argue that scientists are unusually comfortable working with numbers as small as 10^{-3} , but we are not convinced that they are accustomed to explicitly dealing with subjective probabilities of this magnitude. One aid that the analyst can provide is a set of familiar reference events, such as the royal flush example if the expert happens to be a poker player, against which relative likelihood can be compared. Greater assistance is usually provided by enriching the model structure, recognizing that most rare events can be decomposed into a sequence of requisite component events. By modeling that sequence, encoding the conditional probability of each given the occurrence of its predecessors, we enable the expert to assess only probabilities of a readily comprehensible magnitude.

Up to now we have addressed only the problem of how a single scientific expert can be aided in developing a quantitative assessment of planetary contamination risks that is logically consistent with his information and judgment. There is no particular reason to think that two different experts will arrive at the same assessments. In developing a planetary quarantine policy, however, NASA is clearly concerned with the matter of consensus. One would hope that disagreement among experts (as has happened in the assessment of P_g) is traceable to differences in the information available to them and hence resolvable by exchanging information. Such intercourse can be greatly facilitated if assessments are built up from a model structure fine enough to demonstrate which specific aspects of the problem (such as existence of liquid water in the equatorial region) are judged differently. Finally, it is apparent that the relative

expertise of most scientists is not uniform over all considerations relevant to Martian contamination. This has already been implicitly recognized in the use made of the Sagan-Coleman model. The community of experts asked to assess the value of m under various different sterilization procedures is almost totally separate from the community considered most expert in assessing P_g . It seems quite likely that modeling more of the uncertain Martian characteristics that underlie the functional relation between $\text{Prob}\{C\}$ and N would allow further differentiation of expertise.

VI RELATED ISSUES

We have argued that the subjective interpretation is the only meaningful interpretation of the word probability, at least in the context of the planetary contamination problem. Our criticism of the Sagan-Coleman model can be summarized as follows. If equation (1) is viewed as the definition of $\text{Prob}\{C\}$, then this quantity does not represent a probability in any sense that we understand the term, because the assumptions of independence are not consistent with the current state of scientific information. A natural question at this point concerns the spirit of the COSPAR resolution. Rather than actually being concerned with the probability of Martian contamination, its authors might simply have intended to require that $\text{Prob}\{C\}$, as defined by (1), not exceed a specified bound. If this is the case, then current standards and the analytical basis on which they rest are appropriate, assuming that we wish to comply in all good faith with the spirit of the resolution. We would contend, however, that in this event, $\text{Prob}\{C\}$ must be viewed as just one of many possible numerical indices related to the threat of contamination, having no particular significance beyond its being institutionalized through an international agreement.

The great advantage of using probability as the quantitative measure of uncertainty lies in the role that probability assessments play within the broader conceptual framework of decision analysis.⁸ The central idea of this analytical discipline for the treatment of decisions under uncertainty is to separate the issues of value and likelihood. On the one hand we assess (appropriately conditional) probabilities for the uncertain events that impinge on our decision, and on the other we associate value

assessments with each of the outcomes ultimately possible as a consequence of the actions taken. The probabilities and value assessments can then be integrated in a systematic and rational way. An application of this procedure to space project planning is given by Matheson and Roths,⁹ and the same conceptual framework underlies the discussion of Levinthal, Lederberg, and Sagan. If uncertainty is measured in any way other than through probability assessments, then we no longer know how to integrate such measures with value assessments to reveal the preferred course of action. It is our belief that the language of subjective probability, as part of the broader conceptual framework of decision analysis, can be of great value in solving the resource allocation problems that arise at all levels of scientific investigation.

REFERENCES

1. E. C. Levinthal, J. Lederberg, and C. Sagan, "Relationship of Planetary Quarantine to Biological Search Strategy," Life Sciences and Space Research, North-Holland, Amsterdam (1968).
2. C. Sagan and S. Coleman, "Spacecraft Sterilization Standards and Contamination of Mars," Astronautics and Aeronautics, Vol. 3, p. 22 (1965).
3. N. H. Horowitz, R. P. Sharp, and R. W. Davies, "Planetary Quarantine I: The Problem and the Agreements," Science, Vol. 155, p. 1501 (1967).
4. B. C. Murray, M. E. Davies, and P. K. Eckman, "Planetary Quarantine II: Soviet and U.S. Practices and Policies," Science, Vol. 155, p. 1505 (1967).
5. S. Shalkowsky and R. C. Kline, Jr., "Analytical Basis for Planetary Quarantine," Environmental Biology and Medicine, Vol. 1, p. 9 (1971).
6. R. A. Howard, "Decision Analysis: Perspectives on Inference, Decision and Experimentation," IEEE Special Issue on Detection Theory and Application, p. 642 (May 1970).
7. B. de Finetti, "Foresight: Its Logical Laws, Its Subjective Sources," Studies in Subjective Probability, Kyberg and Smokler (eds.), (Wiley, New York, p. 97, 1967).
8. R. A. Howard, "The Foundations of Decision Analysis," IEEE Transactions on Systems Science and Cybernetics, Vol. SSC-4, No. 3 (1968).
9. J. E. Matheson and W. J. Roths, "Decision Analysis of Space Projects: Voyager Mars, presented at the National Symposium, Saturn/Apollo and Beyond, American Astronomical Society (June 11, 1967).

REFERENCES

1. E. C. Levinthal, J. Lederberg, and C. Sagan, "Relationship of Planetary Quarantine to Biological Search Strategy," Life Sciences and Space Research, North-Holland, Amsterdam (1968).
2. C. Sagan and S. Coleman, "Spacecraft Sterilization Standards and Contamination of Mars," Astronautics and Aeronautics, Vol. 3, p. 22 (1965).
3. N. H. Horowitz, R. P. Sharp, and R. W. Davies, "Planetary Quarantine I; The Problem and the Agreements," Science, Vol. 155, p. 1501 (1967).
4. B. C. Murray, M. E. Davies, and P. K. Eckman, "Planetary Quarantine II: Soviet and U.S. Practices and Policies," Science, Vol. 155, p. 1505 (1967)
5. S. Shalkowsky and R. C. Kline, Jr., "Analytical Basis for Planetary Quarantine," Environmental Biology and Medicine, Vol. 1, p. 9 (1971).
- 6. R. A. Howard, "Decision Analysis: Perspectives on Inference, Decision and Experimentation," Proceedings of the IEEE, Vol. 58, p. 632 (May 1970)
7. B. de Finetti, "Foresight: Its Logical Laws, Its Subjective Sources," Studies in Subjective Probability, Kyberg and Smokler (eds.), (Wiley, New York, p. 97, 1967).
8. R. A. Howard, "The Foundations of Decision Analysis," IEEE Transactions on Systems Science and Cybernetics, Vol. SSC-4, No. 3 (1968).
9. J. E. Matheson and W. J. Roths, "Decision Analysis of Space Projects: Voyager Mars, presented at the National Symposium, Saturn/Apollo and Beyond, American Astronomical Society (June 11, 1967).

32. (revised February 1973)